ARL-TR-7638 ● MAR 2016

US Army Research Laboratory

# Autonomous Exploration Using an Information Gain Metric

by Nicholas C Fung, Carlos Nieto-Granda, Jason M Gregory, and John G Rogers

## NOTICES

### Disclaimers

ARL-TR-7638 ● MAR 2016

US Army Research Laboratory

# Autonomous Exploration Using an Information Gain Metric

by Nicholas C Fung, Jason M Gregory, and John G Rogers
*Computational and Information Sciences Directorate, ARL*

Carlos Nieto-Granda
*Institute for Robotics & Intelligent Machines, Georgia Institute of Technology, Atlanta, GA*

| REPORT DOCUMENTATION PAGE | *Form Approved* *OMB No. 0704-0188* |
|---|---|



| 1. REPORT DATE *(DD-MM-YYYY)* | 2. REPORT TYPE | 3. DATES COVERED (From - To) |
|---|---|---|
| March 2016 | Final | 09/2014–09/2015 |

**4. TITLE AND SUBTITLE**
Autonomous Exploration Using an Information Gain Metric

**5a. CONTRACT NUMBER**

**5b. GRANT NUMBER**

**5c. PROGRAM ELEMENT NUMBER**

**6. AUTHOR(S)**
Nicholas C Fung, Carlos Nieto-Granda, Jason M Gregory, and John G Rogers

**5d. PROJECT NUMBER**
R.0013626.6.163.1

**5e. TASK NUMBER**

**5f. WORK UNIT NUMBER**

**7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)**
US Army Research Laboratory
ATTN: RDRL-CII-A
2800 Powder Mill Road
Adelphi, MD 20783-1138

**8. PERFORMING ORGANIZATION REPORT NUMBER**
ARL-TR-7638

**9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)**

**10. SPONSOR/MONITOR'S ACRONYM(S)**

**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**

**12. DISTRIBUTION/AVAILABILITY STATEMENT**
Approved for public release; distribution is unlimited.

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**
Autonomous exploration is a critical capability for robots deployed in real-world situations. Often, these robots are required to navigate environments with little or no prior knowledge in an efficient manner. We investigate an approach to such exploration that selects between candidate frontiers based on the expected information gain. The system uses entropy values of the area map as a measure of the change in overall information and employs established simultaneous localization and mapping (SLAM) techniques. An additional factor considered is the benefit from revisiting previously mapped areas to create loop closures to increase accuracy. We document the approach, including the measure of entropy and loop-closure method. The system is evaluated through simulations and experiments with a live robot. We also evaluate different values for the weighting factors of the utility function, biasing the behavior toward either exploration or relocation. The simulated experiments demonstrate improved efficiency and accuracy in exploring and mapping an unknown environment when compared against a nearest-frontier approach. Results from real-world implementation of the system on robotic platforms, as well as experiments varying the weighting factors of the utility function on the implemented system, are also presented.

**15. SUBJECT TERMS**
information gain, exploration, intelligent behaviors, simultaneous localization and mapping, SLAM, entropy

| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
|---|---|---|---|---|---|
| | | | | | Nicholas C Fung |
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | 19b. TELEPHONE NUMBER (Include area code) |
| Unclassified | Unclassified | Unclassified | UU | 24 | 301-394-3101 |

**Standard Form 298 (Rev. 8/98)**
**Prescribed by ANSI Std. Z39.18**

# Contents

## List of Figures

## 1. Introduction

Robust and efficient autonomous exploration capabilities are a vital component for robot platforms employed in remote and dangerous environments such as providing humanitarian aid in disaster relief scenarios. In such scenarios, there can be limited knowledge of the area and that knowledge may be unreliable. There can also be time constraints, putting an emphasis on the efficiency and accuracy of the robot's capabilities. In the often-used example scenario of exploring a building after an earthquake, the robot must be able to map and navigate around shifted hallways and blocked doors to quickly locate survivors.

This report documents the implementation of an autonomous exploration system designed to quickly and accurately explore and map an unknown area. The system identifies candidate exploration goals using frontiers, then uses a search-based lattice planning algorithm similar to that used by Cohen et al.[1] to generate potential exploration routes. Each route is evaluated based on the predicted entropy change, cost to travel the route, and potential loop closures that can increase the accuracy of the map. Loop closure describes the ability of the robotic system to increase the accuracy of its internal localization based on revisiting a previously explored location. This process reduces the error that can accumulate from inaccuracies in sensor data as the robot constructs the map.

As the robot continues to take sensor measurements and unknown grid cells become known, the entropy of the map is reduced. Because the system prioritizes exploration to areas where the most information can be gained, the rate of exploration is accelerated. In this report, we compare our approach against a baseline nearest proximity frontier selection method.

The *frontier-based* exploration strategy was first described by Yamauchi.[2] In this exploration strategy, frontiers are defined as regions that lie on the boundary between open and unexplored regions on an evidence grid. A robot employing this strategy will make observations at these frontiers, measuring unknown grid cells with its sensors until only open grid cells border obstacle cells and unknown cells are only found on the outside of these obstacles. This strategy is referred to as the "baseline" technique in this report.

Exploration and mapping were linked via an information gain metric from Bourgault et al.[3] by choosing control policies that maximize the information in an occupancy grid and minimize the entropy in the vehicle pose and map features. In Moorehead et al.,[4] map traversability is considered in addition to multiple sources of information, and the total path cost is minimized.

In Sim et al.,[5] the authors explicitly consider optimizing the quality of the map in addition to the speed with which a map is collected by autonomous exploration. This is done by searching for plans that will reenter previously visited locations via a path that will minimize positional uncertainty. In Makarenko et al.,[6] the authors consider the *localizability* of landmarks when deciding on what area to visit next.

Many of these ideas were brought together by Stachniss et al.,[7] who described a system featuring an exploration strategy that maximized occupancy grid information while minimizing map entropy. The system documented in this report uses frontiers as well as predicted locations for loop closures from work documented previously[8] as test locations to evaluate a utility function based on occupancy grid information, map entropy, and path cost. This report is primarily an extension of Stachniss et al.[7] who used modern *pose graph smoothing* simultaneous localization and mapping (SLAM) techniques to predict the effect of planned trajectories on map entropy, including the effect of predicted *loop closure*.

Other work[9,10] has detailed an approach to exploration that uses entropy metrics for path selection in addition to revisiting actions to improve localization. Our technique is most similar to this work; the key difference being that we use the *iSAM2* nonlinear optimization engine to efficiently compute the effect of potential loop closures, which are discovered by executing kinematically feasible trajectories, while the previous work[9,10] used an approximate sparse information filter to compute the value of the information, which can be added through loop closure while following a probabilistic road map. It is not clear from Valencia et al.[9] if their approach is scalable to areas larger than a few rooms, whereas our approach has been evaluated on medium-sized floor plans in this work, and our mapping system on large buildings in Rogers et al.[11]

Here, we document a method of autonomous exploration that uses an information gain metric and SLAM techniques to map an unknown area. We show that this technique improves the efficiency and accuracy of existing techniques and that it can be implemented on widely available robot hardware.

Descriptions of the software components used in this work are presented in Section 2. Exploration experiments in 4 different simulated environments in addition to a real-world experiments are described in Section 3. Finally, conclusions and future plans are described in Section 4.

## 2. System

### 2.1 Exploration Frontiers

The system regulates its autonomous exploration through a series of decisions to select from candidate frontiers and navigation paths. Frontiers are identified as areas that, by some metric, are beneficial for the system to explore. In general, exploration frontiers are locations that the system has identified as areas from which the robot can take measurements and obtain information on unexplored map areas. Such frontiers can be considered candidates for navigation goals, serving to drive an autonomous system. By continually moving to these navigation goals and taking measurements, the system works to explore and map the unknown area.

Exploration frontiers can be described as areas that lie on the boundary between known and unknown space. This type of frontier was initially described by Yamauchi.[2] Maps can be represented as occupancy grids, dividing the floorplan into cells. Each cell is classified as *Clear* if the area is free, *Occupied* if the area is an obstacle, or *Unknown* if the area is unexplored. Frontiers are computed by identifying all occupancy grid cells that are marked Unknown and are immediately adjacent to at least 1 grid cell, which is marked Clear. Groups of these boundary cells are clustered; clusters that consist of a sufficient number of cells are selected as exploration frontiers.

The system chooses an exploration path by evaluating candidate paths to each of the identified exploration frontiers. Candidate paths are evaluated according to the expected information gain from sensor measurements taken along the path and at the candidate frontier. The path that has the greatest reduction in entropy is identified as the best navigation path.

### 2.2 Information Gain Measurement

Entropy is a measurement of uncertainty in the state of a system. A completely unknown environment is at a maximum entropy state, as each grid cell is Unknown and can be said to be either Clear or Occupied with equal probability. As the robot explores the area and takes measurements on these grid cells, the system becomes known and the probability of a measured cell being Clear collapses to 1 (if the cell contains no obstacle) or 0 (if the cell contains an obstacle).

The system endeavors to put the entire area into a low entropy state. That is, it takes measurements of grid cells that are Unknown or have an Occupied probability of 0.5. Therefore, a navigation path that measures and maps a large amount of unexplored environment is desirable. The entropy of a single cell $H_c$ within the

occupancy grid can be calculated according to Eq. 1, where $p(c)$ represents the probability of occupancy of cell $c$. The entropy $H_{og}$ of an occupancy grid $M$ is given by Eq. 2, and is simply a sum across all cells within the grid:

$$H_c(C) = -p(c)\log p(c) - (1-p(c))\log(1-p(c))\,. \tag{1}$$

$$H_{og}(M) = -\sum_{c \in M} p(c)\log p(c) + (1-p(c))\log(1-p(c))\,. \tag{2}$$

As the robot moves through the environment, it collects data through range sensors such as LiDAR. The robot platform was equipped with a Hokuyo UTM-30LX-EW laser range scanner that collects data in a 240°, forward-facing arc. Each individual ray was modeled according to the probabilistic sensor model described in Thrun et al.[12] The model accounts for inherent inaccuracies in beam-based range finders, returning the highest probability at ranges near the ground-truth distance. This distance is calculated through a ray casting operation in the current occupancy grid, stopping at the nearest obstacle. The sensor model is a probability distribution over range $p(r)$, which is instantiated with this distance. The expected information gain $E[I]$ for each ray can be calculated from the sensor model and the cell values resulting from the corresponding ray trace. Equation 3 shows that $E[I]$ is the integration along the ray trace of the probability $p(r)$ that the sensor returns a range value \$r\$ multiplied by the entropy of the cell at this range $H_c(r)$:

$$E[I] = \int_{\mathrm{r}} p(r) H_c(r)\,. \tag{3}$$

The expected information gain is the metric of primary interest when establishing exploration routes.

## 2.3 Mapping

Building a map of the environment is often the ultimate goal of an exploration system. The mapping component of the exploration system, called *OmniMapper*, fuses odometry and laser sensor measurements using square-root smoothing and mapping $\left(\sqrt{SAM}\right)$,[13] implemented in the *GTSAM* library developed at The Georgia Institute of Technology. Specifically, the iSAM2[14] implementation within GTSAM is used as the map optimizer in this work. OmniMapper has been used in prior work to build maps of indoor and outdoor environments[11] and for multirobot mapping.[15–17]

OmniMapper builds a pose graph where nodes are poses along the robot's trajectory and edges are measurements between poses. OmniMapper uses the GTSAM library to solve the pose-graph, which consists of iterative closest point (ICP)[18] scan

matching as well as robot odometry. This enables operation in more general environments, both cluttered and austere. The ICP algorithm used to make relative-pose and loop-closure measurements is canonical scan matching (CSM).[19] This algorithm makes point-to-line measurements to eliminate error caused by limited aperture and handle sparse measurements.

ICP algorithms consist of a 2-phase process, which is repeated until convergence. Two point clouds are given as inputs to the algorithm. For relative-pose measurements, these 2 point clouds come from the robot's current and immediately previous laser scan. For loop-closure measurements, these 2 point clouds are the robot's current and a cloud taken from the same area but from a previous visit. In the first phase of the ICP algorithm, putative point matches are determined by selecting closest matches for each point in the first cloud in the second cloud. In the second phase of the ICP algorithm, these putative point matches form a set of measurements, which is solved in closed form to find the relative pose that best brings these 2 clouds into alignment. This transformation is then applied to the first set of input points. A new set of putative point correspondences are selected using this updated point cloud, and the algorithm repeats until convergence, when the set of corresponding points is unchanged.

## 2.4 Loop Closures

In addition to the uncertainty related to the unexplored portion of the environment, there exists uncertainty in the instrumentation and interpretation of data within the robotic platform. Many robotic mapping systems rely on sensor measurements such as odometry readouts, inertial measurement units, and LiDAR scanners to determine their own position and interpret the environmental data measurements. Without proper localization, constructed maps of the environment will contain errors, such as slanted hallways.

The system incorporates the ability to relocate itself within its internally constructed map by interpreting sensor data from the laser range scanner and the odometry sensors. By revisiting areas that have already been mapped, the system can resolve conflicts between observed and expected sensor readings, also known as loop closures. The ICP procedure described in Section 2.3 is used to determine the relative pose between the robot's current and previous position when revisiting a previously mapped location.

As the robot builds a map of an environment, the robot's pose uncertainty will increase unless loop-closure measurements are made. This uncertainty can be quantified by computing the entropy of the robot's a posteriori pose estimate. The robot's pose history along its trajectory is captured by the mapping operation

detailed in Section 2.3. The marginal distribution can be efficiently computed using the iSAM2.[14] The entropy of the most recent trajectory element is a measurement of the robot's pose uncertainty in the map, is computed with Eq. 4. In this expression, the marginal covariance on pose *i* is given by $\Sigma_{ii}$, and $H_p$ is the differential entropy of the robot's pose:

$$H_p = \frac{1}{2}\log (2\pi e)^3 \left|\sum_{ii}\right| . \tag{4}$$

As the robot explores its environment, it considers a small set of potential destinations before selecting 1 destination as the next exploration goal. The effect of selecting each potential destination as the next goal with respect to the trajectory entropy expression in Eq. 4 can be approximated. For each potential destination, a plan is computed, which will reach that destination with a search-based lattice planner (Search-based Planning Laboratory [SBPL]).[1] The effects on the mapper trajectory of following this plan are hypothesized by appending trajectory pose variables along the plan. Edges are added between adjacent trajectory pose variables, and loop-closure edges are added whenever the robot's planned trajectory gets close to a previously visited area. When these loop-closure edges are added to a hypothesized trajectory based on a plan, the entropy of the mapper can be reduced.

### 2.5 Utility Function

The system calculates the utility of a potential path using Eq. 5. $\Delta H_M(p)$ is the change in map entropy, Eq. 2, as a result of traveling path *p* while taking sensor measurements. $\Delta H_X(p)$ is the change in entropy from the pose change, Eq. 4, over path *p*, weighted by constant *k*. This constant can be adjusted to increase or decrease favor of loop closures. $C(p)$ is the cost to travel along the path, which is a function of the robot:

$$U(p) = \frac{\Delta H_M(p) + k\Delta H_X(p)}{C(p)} . \tag{5}$$

## 3. Experimentation

The exploration system was evaluated in both simulated and real-world environments. Experimentation in the simulation world was performed in a number of different environments designed to evaluate the system under a variety of floorplans.

## 3.1 Simulation

The information-based system was evaluated against a baseline system using a proximity frontier selection method in 2 experiments. The first experiment was meant to evaluate the efficiency of the algorithms by recording the occupancy grid entropy of the map as a function of time. Four maps of varying environments, depicted in Fig. 3, were used to enforce robust performance of both algorithms. The second experiment evaluated the ability of the system to form an accurate map. This experiment introduced noise to the simulated sensor measurements of the robot, causing mapping errors and making it more difficult to build a coherent, accurate map. In this experiment, forming loop closures would be vital in creating an accurate map.

In real-world operation, the system must be flexible in its ability to expand the map beyond the current limitations. However, the map is stored as a rectangular grid in memory. As the map expands, the overall entropy of the system increases because the system must allocate many new, unknown cells to maintain the rectangular map. This increases in entropy is contrary to the systematic goal of entropy minimization. In practice, this does not affect the system calculations because the calculation for expected entropy change does not include the addition of unknown cells. However for evaluation purposes, the simulation allocates a large number of unknown cells to the entropy calculation, padding the map to a size that is large enough to cover the map. By using this padding method, the entropy value should have an overall decreasing trend across each trial.

## 3.2 Robotic Platform

The exploration system was implemented on 2 different robotic platforms: the PackBot designed by iRobot and the Jackal designed by Clearpath Robotics. The PackBot, shown in Fig. 1, is a man-portable robot system. The robot was equipped with additional computing hardware to increase the capabilities of the platform. Similarly, the Jackal is a wheeled, man-portable robot system. Both robots were equipped with a Hokuyo UTM-30LX-EW scanning laser range finder with a motor controller designed to continually nod the sensor. This sensor provides 3-dimensional point cloud data to the platform.

**Fig. 1 PackBot used to perform the experiments**

Experimentation on the PackBot platform has primarily been carried out on an informal basis intended to show "proof of concept" in integrating the decision-making system on to a real-world robotic system. As shown in Fig. 2, the robot was used to explore and map the second floor of a building located in a military and rescue training facility. The Jackal platform was used to collected experimental data designed to evaluate the exploration system. Figure 3 shows the simulation environments.

**Fig. 2 PackBot exploring an abandoned pediatric mental hospital, currently part of a military and rescue training facility**

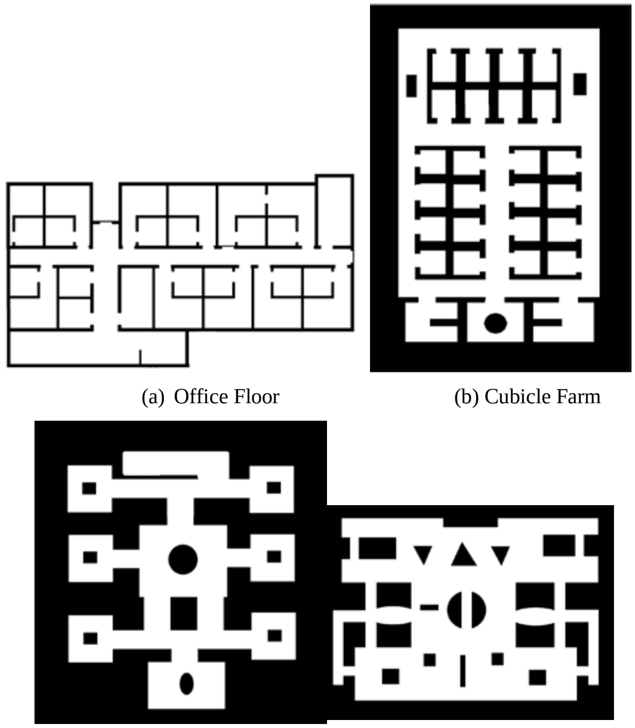

(a) Office Floor (b) Cubicle Farm

(c) Underground Lair (d) Reactor Core

**Fig. 3 Simulation environments used to perform the exploration experiments**

## 3.3 Simulation Results

Results from the first simulation experiment meant to evaluate efficiency can be seen in Fig. 4. The figure shows the average of 5 trial runs for each algorithm per map. The information-based approach was significantly faster than the baseline approach in 3 of the 4 maps, attaining a lower occupancy grid entropy state over the entire time period. For the fourth map, the information-based approach attained a small advantage, although both algorithms reach a similar steady entropy state quickly.

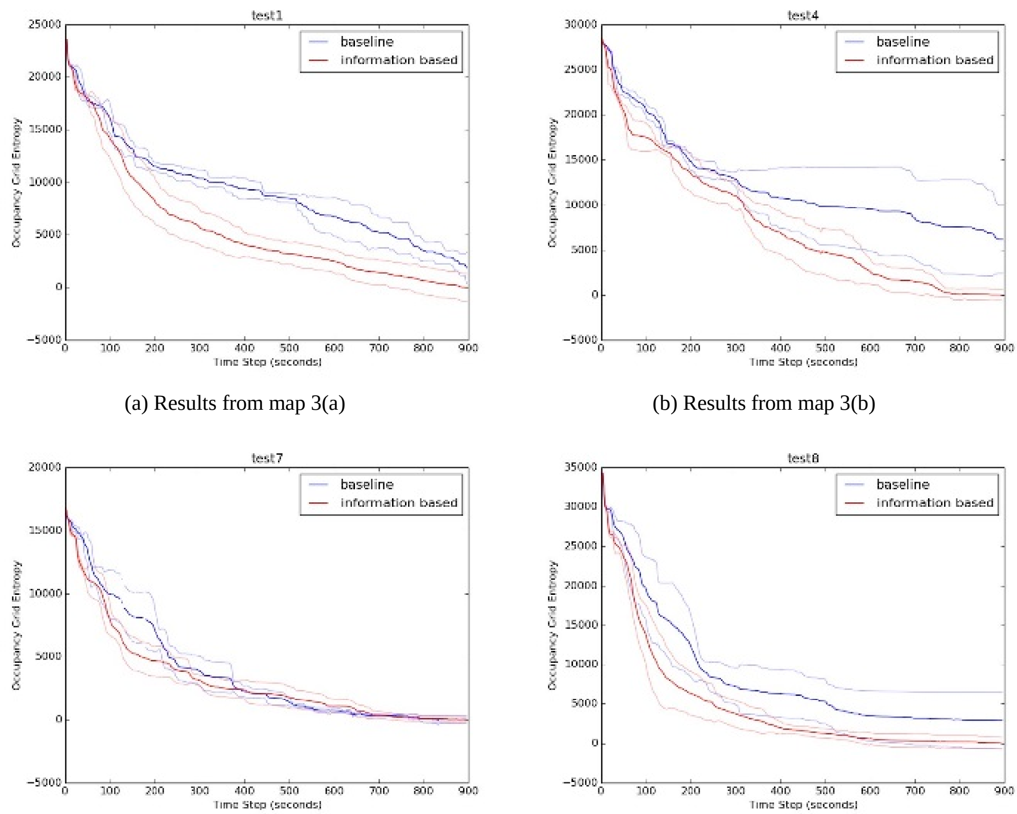


(a) Results from map 3(a)

(b) Results from map 3(b)

(c) Results from map 3(c)

(d) Results from map 3(d)

**Fig. 4 Occupancy grid entropy as a function of time averaged over 5 runs for each of the map. In all graphs, red represents the information-based algorithm and blue represents the baseline algorithm. A 1σ error envelope is drawn around each line. This figure is best viewed in color.**

Results from the second experiment meant to evaluate the accuracy of the algorithms can be seen in Fig. 5. Each algorithm was run for 42 trials on the map shown in Fig. 3a. Figure 5 depicts the 3 best maps for each algorithm as selected by a third party. The information-based exploration was set to favor loop closure. This experiment was run with large sensor noise variance, which caused mapping failure (duplicated structures, closed-off spaces) that would prevent successful navigation. The information-based exploration maps were generally of better quality and covered most of the environment, as seen in Fig. 5a. Baseline exploration maps encountered mapping failure before covering more of the environment than is seen in Fig. 5b.

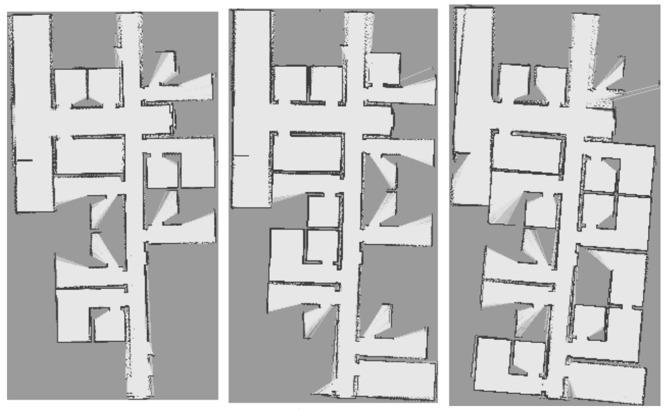

(a) Three best information-based runs

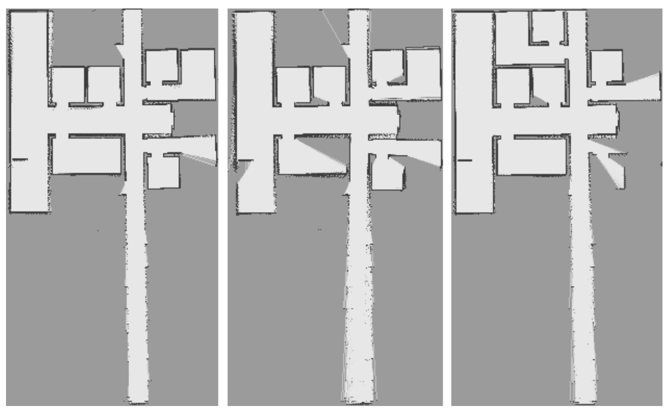

(b) Three best baseline runs

**Fig. 5 Best 3 maps from 42 runs each of baseline exploration and information-based exploration, terminated after 10 min of runtime**

In addition to the 2 formal experiments, the trajectory of the baseline algorithm was compared against the information-based algorithm with various weights given loop-closure favorability ($k$ from Eq. 5). Some representative runs can be seen in Fig. 6. Each of these runs was terminated after 5 min to increase readability of the robot trajectories in the figure. Figure 6a shows that the baseline algorithm explores a compact area, as it methodically proceeds to the nearest frontier. Figure 6b shows the information based algorithm with equal loop-closure weighting, a typical setting for this approach. In this run, the robot explores a medium amount while trying to form some loop closures. In Fig. 6c, the loop-closure weighting is set to 0 so that the robot always chooses the most informative frontier, which is typical far away from the previous pose, as it will uncover a lot of information. Figure 6d depicts the information-based approach when the loop-closure weight is larger than the map information weight. This configuration results in the robot returning to previously explored areas to form loop closures.

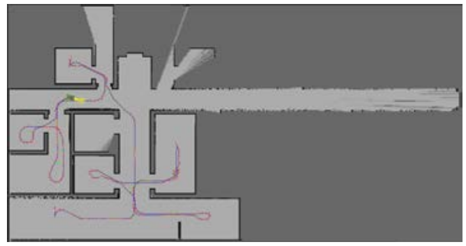

(a) Baseline

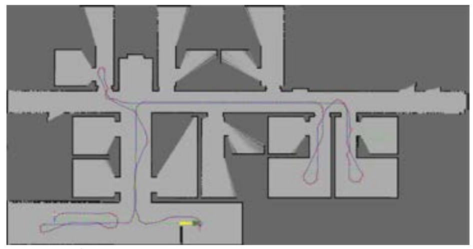

(b) Information based. Loop closure given equal weighting to map information gathering.

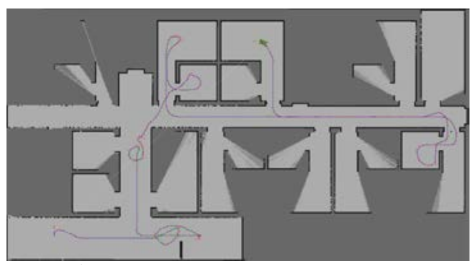

(c) Information based. Loop closure weight 0.

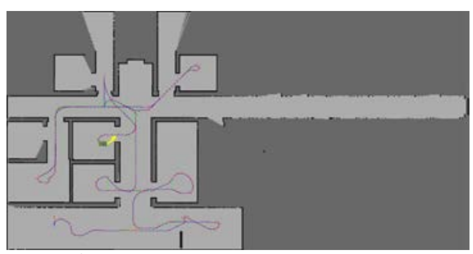

(d) Information based. Loop closure given higher weighting than map information gathering.

**Fig. 6 Example trajectories from various settings of loop-closure weight parameters as well as baseline system after 5 min of execution**

## 3.4 Implementation Results

The exploration system was implemented on 2 different robotic platforms. A PackBot was used to map a hospital building at a military and rescue training facility. The results can be seen in Fig. 7. The map shows good overall coverage of a complex and sizeable area. The exploration system continually prioritized frontiers on the fringe of the explored area and mapped the majority of the side rooms. This system was also evaluated on other floors in this training facility as well as in other test environments with similar results.

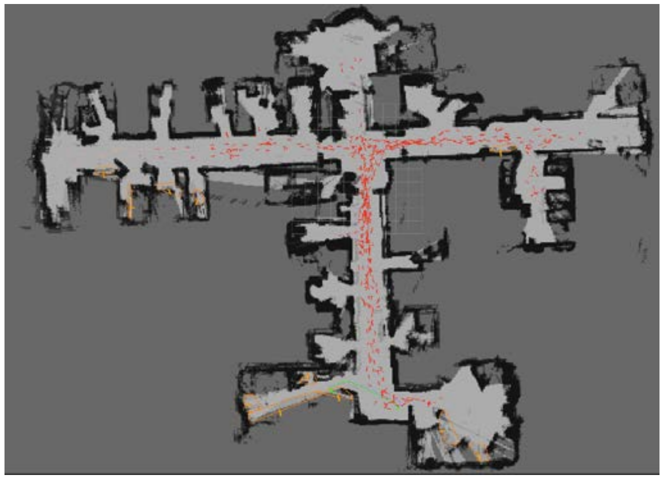

**Fig. 7 Mapping of the second floor "hospital" building (abandoned children's mental hospital), the urban rescue training facility**

The system was also implemented on a Clearpath Robotics Jackal robot. This robot was used to explore and map a training structure resembling a police station. For this experiment, the scaling factor of the graph entropy was varied. The results can be seen in Fig. 8. This experimentation demonstrates the ability for the algorithm to be implemented on a second platform from the PackBot. In addition, the data collected over a variety of graph entropy scale values demonstrate the behavior of the algorithm when loop closures are increased or decreased in priority. Figure 8 shows that smaller graph entropy scale values resulted in the exploration algorithm more quickly reducing the map entropy versus the baseline exploration algorithm. However, when the scale value is increased to 1000, the baseline performs better. This results from over adjusting the scale value, making the system strongly value forming loop closures. As a result, the robot revisits mapped areas at a cost to quickly exploring the area.

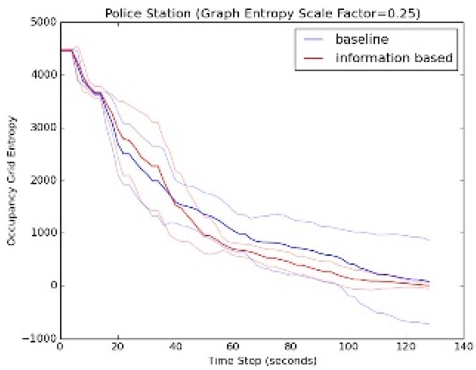


(a) Graph Entropy Scaling=0.25

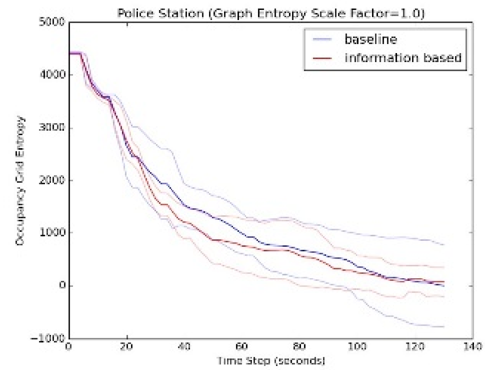


(b) Graph Entropy Scaling=1.0

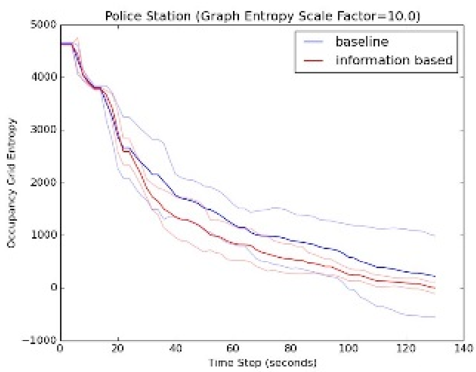


(c) Graph Entropy Scaling=10.0

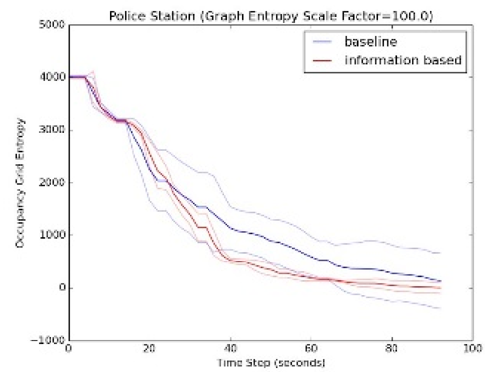


(d) Graph Entropy Scaling=100.0

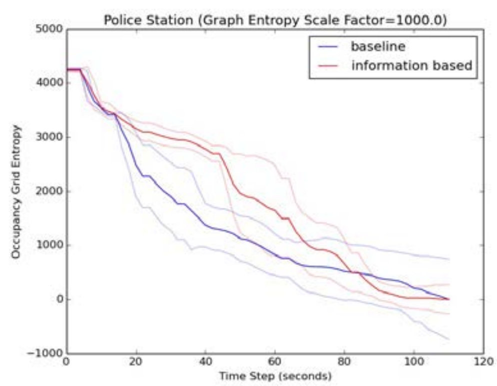


(e) Graph Entropy Scaling=1000.0

**Fig. 8 Entropy values while mapping a police station building**

## 4. Conclusions

The information-based exploration system performed at a more efficient and accurate level than the baseline proximity frontier selection method. The system can be extended to a number of different implementations such as multiple robot agents or unmanned aerial vehicles without regard to the details of control or sensing. The real-world mapping depicted in Fig. 7 showed that the system could be implemented on a currently deployable platform constrained by energy and computing power. The system incorporates established research in the area of information-based exploration with pose graph smoothing SLAM techniques to produce an effective exploration technique. The system is able to create coherent maps efficiently and use loop closures to increase accuracy.

## 5. References


1. Cohen BJ, Chitta S, Likhachev M. Search-based planning for manipulation with motion primitives. In: IEEE 2010. Proceedings of 2010 IEEE International Conference on Robotics and Automation (ICRA); 2010 May 3–8; Anchorage, AK. p. 2902–2908.

2. Yamauchi B. A frontier-based approach for autonomous exploration. In: CIRA 1997. Proceedings of 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation; 1997 Jul 10–11; Budapest, Hungary. p. 146–151.

3. Bourgault F, Makarenko A, Williams S, Grocholsky B, Durrant-Whyte H. Information based adaptive robotic exploration. In: IROS 2002. Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems, Vol. 1; 2002 Sep 30–Oct 4; Lausanne, Switzerland. p. 540–545.

4. Moorehead S, Simmons R, Whittaker W. Autonomous exploration using multiple sources of information. In: ICRA 2001. Proceedings of IEEE International Conference on Robotics and Automation, Vol. 3; 2001 May 21–26; Seoul, Korea. 3098–3103.

5. Sim R, Dudek G, Roy N. Online control policy optimization for minimizing map uncertainty during exploration. In: ICRA 2004. Proceedings of 2004 IEEE International Conference on Robotics and Automation, Vol. 2; 2004 Apr 26–May 1; New Orleans, LA. p. 1758–1763.

6. Makarenko AA, Williams SB, Bourgault F, Durrant-Whyte HF. An experiment in integrated exploration. In: IROS 2002. Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems, Vol. 1; 2002 Sep 30–Oct 4; Lausanne, Switzerland. p. 534–539.

7. Stachniss C, Grisetti G, Burgard W. Information gain-based exploration using rao-blackwellized particle filters. In: RSS 2005. Proceedings of Robotics: Science and Systems (RSS); 2005 Jun 8–11; Cambridge, MA. p. 65–72.

8. Stachniss C, Hahnel D, Burgard W. Exploration with active loop-closing for fastslam. In IROS 2004. Proceedings of 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vol. 2; 2004 Sep 28–Oct 2; Sendai, Japan. p. 1505–1510.

9. Valencia R, Miró JV, Dissanayake G, Andrade-Cetto J. Active pose slam. In: IROS 2012. Proceedings of 2012 IEEE/RSJ International Conference on

Intelligent Robots and Systems (IROS); 2012 Oct 7–12; Coimbra, Portugal. p. 1885–1891.

10. Vallvé J, Andrade-Cetto J. Dense entropy decrease estimation for mobile robot exploration. In. ICRA 2014. Proceedings of 2014 IEEE International Conference on Robotics and Automation (ICRA); 2014 May 31–Jun 7; Hong Kong, China. p. 6083–6089.

11. Rogers III JG, Young S, Gregory J, Nieto-Granda C, Christensen HI. Robot mapping in large-scale mixed indoor and outdoor environments. In: Proc. SPIE Defense and Security; 2013 Apr 29–May 3; Baltimore, MD.

12. Thrun S, Burgard W, Fox D. Probabilistic robotics (intelligent robotics and autonomous agents). Cambridge (MA): The MIT Press; 2005.

13. Dellaert F. Square root SAM: Simultaneous localization and mapping via square root information smoothing. In: RSS 2005. Proceedings of Robotics: Science and Systems (RSS); 2005 Jun 8–11; Cambridge, MA. p. 1181–1203.

14. Kaess M, Ranganathan A, Dellaert F. Fast incremental square root information smoothing. In: IJCAI 2007. Proceedings of International Joint Conference on Artificial Intelligence (IJCAI); 2007 Jan 6–12; Hyderabad, India.

15. Rogers JG, Nieto-Granda C, Christensen HI. Collaboration strategies for multi-robot exploration and mapping. In: ISER 2012. Proceedings of International Symposium on Experimental Robotics (ISER); 2012 Jul 21–25; Berlin, Germany.

16. Rogers III JG, Nieto-Granda C, Christensen HI. Collaboration strategies for multi-robot exploration and mapping. In: Desai JP, Dudek G, Khatib O, Kumar V, editors. Experimental robotics: the 13th international symposium on experimental robotics, Part III; 2013. Cham (Switzerland): Springer International Publishing. p. 231–243.

17. Rogers J, Paluri M, Cunningham A, Christensen HI, Michael N, Kumar V, Ma J, Matthies L. Distributed autonomous mapping of indoor environments. In: Proc. SPIE Defense and Security; 2012 Apr 23–27; Baltimore, MD.

18. Besl P, McKay N. A method for registration of 3-D shapes. IEEE Trans on Pattern Analysis and Machine Intel. 1992;14(2):239–256.

19. Censi A. An ICP variant using a point-to-line metric. In: ICRA 2008. Proceedings of IEEE International Conference on Robotics and Automation; 2008 May 19–23; Pasadena, CA. p. 19–25.

| | |
|---|---|
| 1 (PDF) | DEFENSE TECH INFO CTR<br>DTIC OCA |
| 2 (PDF) | US ARMY RSRCH LAB<br>IMAL HRA MAIL & RECORDS MGMT<br>RDRL CIO LL TECHL LIB |
| 1 (PDF) | GOVT PRINTG OFC<br>A MALHOTRA |
| 3 (PDF) | US ARMY RSRCH LAB<br>RDRL CII A<br>N FUNG<br>J GREGORY<br>J ROGERS |